

Cross-layer Energy-Efficiency Optimization of Packet Based Wireless MIMO Communication Systems

Christian Senning · Georgios Karakonstantis · Andreas Burg

Received: date / Accepted: date

Abstract Energy in today's short-range wireless communication is mostly spent on the analog- and digital hardware rather than on radiated power. Hence, purely information-theoretic considerations fail to achieve the lowest energy per information bit and the optimization process must carefully consider the overall transceiver. In this paper, we propose to perform cross-layer optimization, based on an energy-aware rate adaptation scheme combined with a physical layer that is able to properly adjust its processing effort to the data rate and the channel conditions to minimize the energy consumption per information bit. This energy proportional behavior is enabled by extending the classical system modes with additional configuration parameters at the various layers. Fine grained models of the power consumption of the hardware are developed to provide awareness of the physical layer capabilities to the medium access control layer. The joint application of the proposed energy-aware rate adaptation and modifications to the physical layer of an IEEE 802.11n system, improves energy-efficiency (averaged over many noise and channel realizations) in all considered scenarios by up to 44%.

Keywords energy-efficiency · MIMO communication · cross-layer optimization

1 Introduction

Mobile communication – anytime, anywhere access to data and communication services – has been continuously increasing since the operation of the first cellular phone system. This growth is combined with an increasing demand by consumers for small and multifunctional products. Such products must be able to transmit wirelessly not only voice, but also color pictures and video, as well as to provide access to complex applications over the worldwide web [9]. Satisfying such a growing demand for wireless communications has been achieved by advances in information theory, combined with the ability to manufacture high throughput communication circuits. Unfortunately, concerns are being raised that this scenario cannot continue forever. Specifically, as more and more transistors are packed onto a single chip to support the associated demand for high performance communications, more of them toggle in “active” mode or leak during “sleep” periods, resulting in a substantial increase of on-chip power dissipation [22]. Such increased power dissipation in combination with the slow improvements in battery capacity limit the battery run-time of portable devices and prevent them from adequately meeting user expectations.

The classical optimization for high-performance in terms of throughput or error rate [14, 26] performed by goodput-guided rate adaptation (RA) schemes may not be able to close the gap between supplied and required energy. This problem may be attributed to the fact that either the best error-rate might not always be needed, or the highest throughput might exceed application requirements. In particular the peak data rates in the widely used IEEE 802.11n standard (600 Mbps) significantly surmount the requirements for high-quality audio transmissions (192 kbps) or high-definition video-

Christian Senning, Georgios Karakonstantis, Andreas Burg
EPFL-STI-IEL-TCL, Station 11, ELG 011,
Tel.: +41 21 693 6924
Fax: +41 21 693 2687
E-mail: {christian.senning, georgios.karakonstantis, andreas.burg}@epfl.ch}
This paper extends the work published in
ICASSP [32]

on-demand services (e.g., YouTube.com limits the maximum rate of its content to 5 Mbps). This means that we have reached a point where the peak data rates of the wireless communication system is no longer always a limiting factor. This can be argued especially for scenarios where the cumulative data rate requested by all users in the same wireless local area network (WLAN) channel is far less than the system capacity, as in isolated networks with small number of users (e.g., networks within detached houses or small offices). Therefore, we show that there are many opportunities which can be exploited by alternative RA schemes to enhance the battery run-time of portable devices.

Some works have proposed energy-efficient RA schemes by mainly focusing on the minimization of the radiated power only. However, such optimization may not be so beneficial from an energy-efficiency point of view for today's networks for short range wireless connectivity, which usually spend most of the power in the signal processing of the receiver [13]. Recent studies shift the throughput objective of RA schemes to energy-efficiency optimization of the overall system [13, 23]. Although, such works improved the energy-efficiency of wireless systems, they have not yet revealed all the potential gains, since i) either they focus solely on the RA, neglecting possible tuning knobs on other layers or ii) their knowledge of the physical (PHY) layer is based on high-level information theoretical energy models, ignoring the inter dependencies across all layers and modules or iii) they investigate a given, fixed PHY layer treated as black-box ignoring the potential to adjust processing energy according to data rate or channel conditions (energy proportionality).

Contributions: In this paper, we propose to enhance the energy-efficiency by a cross-layer approach, based on energy-aware RA schemes combined with an *energy proportional* PHY layer design i.e., a receiver that is able to scale its effort (i.e., its energy consumption) according to the varying complexity to successfully recover the transmitted bits under all given operating conditions. Specifically, our contributions can be summarized as follows:

- An energy-guided RA scheme is developed and the achievable gains in terms of energy spend per successfully received information bit are analyzed. Our results indicate that such a RA scheme can result in significant improvements in energy-efficiency, while still satisfying the data rate of most today's applications.
- A compromise between the energy-guided RA and the classical goodput-guided RA is presented for applications with high data rate requirements.
- Both proposed RA schemes are enabled through a fine-grained energy model of the PHY layer that allows to capture many inter-dependencies across various settings throughout the operation of a packet based wireless system.
- The available choices of classical RA, assumed by most existing works, are extended by exploiting circuit level techniques combined with algorithmic effort scaling. This allows to truly realize an energy proportional PHY layer and maximize the energy-efficiency gains compared to the classical throughput-guided RA that does not exploit any of the proposed PHY layer modifications.
- A case study based on an IEEE 802.11n compliant PHY is shown for fixed and varying packet lengths.

Outline: Section 2 discusses prior publications for energy-efficient wireless communication systems and corresponding energy models. Section 3 elaborates on the background regarding a typical frame-based communication system and Section 4 presents the proposed approach. The required energy model is developed in Section 5, while the applied techniques for enabling an energy proportional PHY layer are described in Section 6. The proposed optimizations are applied to a typical IEEE 802.11n environment and the results are presented in Section 7. Finally, conclusions are drawn in Section 8.

2 State of the Art

The traditional research has been focusing on optimizing the transmit power based on information-theoretic energy-efficiency metrics [8, 16]. However, the proposed solutions start from a model in which the device power consumption is essentially given by the radiated energy [24, 29]. Unfortunately, this assumption rarely holds in practice as for most battery-operated systems used for wireless connectivity (e.g., WLAN) the radiated power is only a small part of its total power consumption. Later studies have recognized this fact and have started to take active power consumed by the hardware into account. While this defies the straightforward application of information theoretic tools and principles, it also provides more opportunities for energy savings [38].

The hardware power has been accounted for in some early studies, e.g., [1]. However, the point-to-point streaming-system under consideration in [1] that only considers a physical layer, the corresponding traffic pattern, and the associated tuning knobs has little in common with today's omnipresent short-range wireless networks (e.g., IEEE 802.11n or IEEE 802.11ac). Based on this observation recent works have been focusing on

providing strategies to improve energy-efficiency at the medium-access-control (MAC) layer [27, 36]. The latest studies on the component level consider different algorithm choices under the aspect of energy-efficiency, spanning both the algorithm and the architectural layer [15].

Later publications introduced the concept of energy-aware baseband processing, where scalability is put to service to adjust the processing to the dynamically changing environment, leading to even better overall energy-efficiency [17, 18, 23]. In this context, a recent work by Intel tried to provide realistic power measurements of the baseband for the popular IEEE 802.11n standard claiming to be the first ever doing it, indicating a shift in design objectives from throughput oriented communications to energy-efficiency oriented communication [13]. In particular, [38] showed that the baseband processing of software-defined radios is most of the time underutilized and states that this underutilization can be used for power reduction (resulting in the desired energy proportional behavior).

In order to exploit such underutilization, the MAC layer has to be aware of the energy consumption of its associated PHY layer. For this purpose, many models for the energy consumption of wireless systems have been proposed in the literature. Some models focus exclusively on the transmit power [16] and do not capture the main energy-drains, while others jointly model the baseband processing of the transmitter and the receiver [3, 35]. These models then also result in optimization of the grid-powered access points instead of focusing on the more critical battery operated devices. Further works rely on measurements of entire chip sets or network interface cards [6, 13, 21, 34] that do not provide insight into the energy consumption of the PHY layer itself.

Although existing studies have indicated the need for improving energy-efficiency, there is still a need for true cross-layer optimization, enabled by an energy-guided RA that is aware of the energy consumption of an associated energy proportional PHY layer, through accurate models.

3 Background

Fig. 1 illustrates a point-to-point link of a frame-based wireless communication system, which is typical for short-range WLAN networks. The essential components for our consideration are a grid powered, central access point (AP) and one or more battery operated devices that are referred to as *clients*. In our scenario, all clients receive application data (e.g., audio or video) transmitted from the AP over a wireless channel.

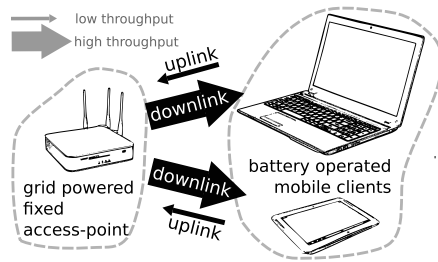


Fig. 1: WLAN setup with high traffic load from access point to battery operated mobile clients.

Both the AP and the client, capable of transmitting and receiving multiple data streams using multiple antennas, consist of a PHY layer that is responsible for the actual transmission and reception of the data as well as a MAC layer that controls the PHY layer. The PHY layer is composed of various hardware modules. In transmit mode, these modules encode the application data, modulate them, convert them into analog radio frequency (RF) signals, and finally radiate them over the available antenna(s). On the other side of the point-to-point link, the PHY layer operates in receive mode and its modules apply the opposite functions converting the received RF signals into digital binary data, demodulating and decoding them as well as finally extracting the application data.

We assume a system that operates in time-division-duplex (TDD) and time-division-multiple-access (TDMA) mode and we consider a typical asymmetric scenario in which data is mostly downloaded from the AP to the client. Hence, the client is most of the time in receive mode. However, there are also some short time-slots during which each client gets into transmit mode for sending acknowledgment (ACK) frames to the AP.

3.1 Packet Structure

As illustrated in Fig. 2, the transmission sequence begins with a frame start waiting period. The duration of this period depends on proper sleep time prediction based on the power save poll MAC protocol. The actual frame starts with a training sequence used for frame start detection, initial frequency offset estimation, and channel estimation. The initial channel estimate is then used to detect the frame header containing information about the subsequent data payload (e.g., the payload length in number of bytes and the modulation and coding scheme (MCS), which together determine the frame length in number of symbols). If the number of spatial streams differs between header and payload, an additional training sequence is required to enable the esti-

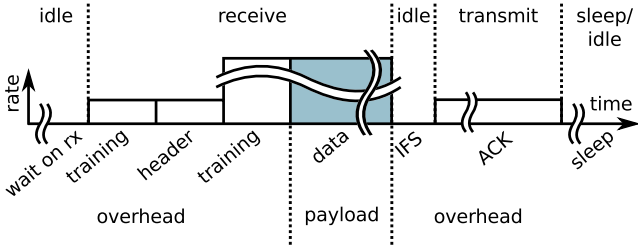


Fig. 2: Simplified illustration of the IEEE 802.11n transmission protocol.

mation of the multiple-input multiple-output (MIMO) channel required later for detection of the data payload. The final part of the received frame is the data payload itself. Within an inter-frame spacing (IFS), the receiver has to decide if the data payload was correctly extracted and has to start transmitting an ACK frame. After sending the ACK frame, the PHY layer goes into sleep mode until he is again triggered by the MAC layer for the next frame reception.

3.2 Classical Rate Adaptation

A key feature of most communication systems for short-range wireless connectivity, as described above, is their ability to support a variety of transmission schemes to adjust the rate to the current channel conditions. For instance, the modulation and code rate can be selected along with other parameters such as the number of spatial streams N_{SS} , which together define the MCS, encoded in the header field of a packet. The transmission mode and the values of the associated parameters are selected by the MAC layer of the transmitter, based on channel state information¹. In today's systems, the main objective of commercially available goodput-guided (GG) RA schemes is the selection of the most appropriate system mode ν under given channel conditions for the maximization of the goodput. The system mode is selected from a set Ω_{GG} that contains all possible transmission modes. Such transmission modes determine the employed MCS and the frame length L in bits. The objective of GG RA is therefore given as follows

$$\nu_{GG} = \arg \max_{\nu} \{ (1 - P_e(\nu)) \Phi(\nu) \}, \quad (1)$$

where $\Phi(\nu)$ is the throughput and $P_e(\nu)$ is the probability of a packet error. The GG RA therefore requires to estimate the error rate at run-time based on the current

¹ While data reception relies on accurate up-to-date channel state information, we have observed that channel characteristics relevant for RA remain stable over a long time.

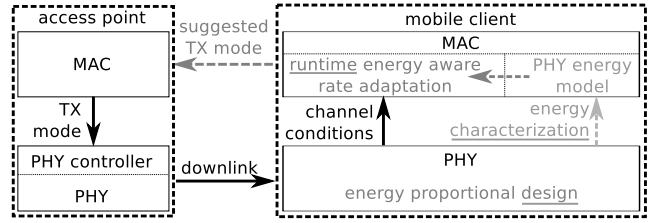


Fig. 3: Proposed approach for improved energy-efficiency of the mobile client.

channel conditions. For this, a variety of known techniques for accurately estimating $P_e(\nu)$ can be applied as discussed for example in [18].

The choice of ν under such a GG RA actually determines the operating mode of both the transmitter and receiver. Due to differences in the (de-)modulation and (de-)coding process as well as differences in the packet structure and duration this can also lead to differences in the required signal processing effort. Hence, the MCS selection by the RA clearly has an impact on system power consumption on both sides of the wireless link.

4 Proposed Cross-Layer Approach for Improved Energy-Efficiency

As we discussed, the client in modern wireless links is in many cases a battery operated device that operates most of the time in receive mode. The amount of data transmitted to the client is usually independent of the available peak data rates but given by use patterns and applications. Therefore, the main goal of our cross-layer energy-efficiency optimization is to reduce the energy consumption per successfully received information bit at the receiver to maximize the number of bits that can be received on a single battery charge.

An illustration of the overall proposed approach is shown in Fig. 3. The medium access control (MAC) layer at the AP configures its PHY layer and forwards data to be transmitted. A packet containing this data is then sent over a wireless channel to the mobile client. The PHY layer of the client tries to recover the transmitted data based on the received signal. The energy consumption of the PHY layer of the client thereby depends on the transmission mode and the “difficulty” to recover the data under given channel conditions. To account for this dependency when choosing a mode, the energy consumption for different configurations are modeled in the MAC layer of the client. Based on a prediction of this energy model as well as on channel conditions reported from the PHY layer, a suitable transmission mode and corresponding receiver configurations

are suggested to the AP for future transmissions. Such a RA is discussed in the next section of this publication.

4.1 Energy-Guided Rate Adaptation

In our approach, we depart from the classical GG RA and propose an energy-guided (EG) RA which aims at reducing the energy consumption of the receiver. Specifically, the method selects the system mode according to

$$\nu_{EG} = \arg \min_{\nu} \{\eta(\nu)\}, \quad (2)$$

with $\nu \in \Omega_{EG}$, where Ω_{EG} corresponds to the set of all meaningful system modes that include the modes in Ω_{GG} as well as receiver specific settings. The function $\eta(\nu)$ represents the energy consumption of the battery operated client per received information bit. Note that, at this point no attention is specifically paid towards the impact on goodput or error rate, except for the implicit dependency of $\eta(\nu)$ on the data rate $\Phi(\nu)$ and the probability of a packet error $P_e(\nu)$. As we will elaborate on, this dependency still provides a small bias toward the higher rate modes, but also respects the availability of modes associated with lower rates for the benefit of energy-efficiency.

In contrast to the classical RA performed at the transmitter (based on channel quality feedback), energy-guided RA has to be performed at the battery operated client. This is because the AP can usually not accurately estimate or model the power consumption of a third-party battery operated device. To this end, we assume a setup in which MCSs are suggested from the client to the MAC layer of the AP² for the next data frame (assuming a reasonably static channel), during the ACK frames.

For occasions, when the goodput achieved by the EG RA does not satisfy the required data rate of an application, we propose a goodput-aware energy-guided (GAEG) RA that provides a compromise between the goodput achieved by the GG RA and the energy-efficiency provided by the EG RA. Such a compromise is achieved by choosing the mode ν_{GAEG} with the best goodput and an energy-efficiency that is upper bounded by the factor $k > 1$ times the energy per bit of ν_{EG} based on

$$\nu_{GAEG} = \arg \max_{\nu} \{(1 - P_e(\nu))\Phi(\nu) \mid \eta(\nu) < k\eta(\nu_{EG})\}. \quad (3)$$

In this approach the MAC layer is able to trade energy consumption versus throughput at run time and therefore allows adjustment to the application requirements or for systems operated close to the capacity bound, by selecting an appropriate value of the factor k .

² Important standards such as IEEE 802.11n already include this possibility

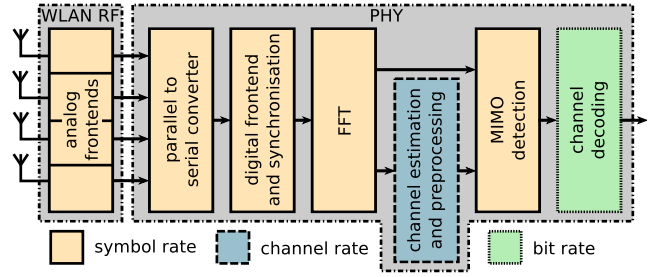


Fig. 4: Simplified PHY layer block diagram with the processing rate of the modules highlighted.

4.2 Augmenting System Modes

In addition to shifting the RA to the client side, the proposed RA in (2) does not only select an MCS and the packet length L , but also other configuration parameters of the client. Such an *extended set* Ω_{GG} of configuration parameters may include: i) algorithmic choices, such as different MIMO detection algorithms, ii) number of active receive chains, and iii) applied circuit level techniques, like clock-gating or dynamic voltage and frequency scaling (DVFS). Due to these extended system modes, the PHY layer can achieve a better energy proportional behavior, compared to the fixed PHY layers that were used in prior works. With these extensions, the RA now has more selection choices, which extend the Pareto-frontier in terms of maximizing the achievable energy-efficiency gains.

5 Receiver Energy Model

To perform the mode selection at run time according to (2), the MAC layer of the client needs to estimate the power consumption of the modules composing the PHY layer for each possible system mode ν . To enable such fine grained energy awareness, we model the energy consumption of the PHY layer by partitioning the contribution of each module based on their participation in the different phases of the transmission protocol shown in Fig. 2.

In order to understand the energy consumption of the PHY and MAC layers, a detailed energy model is needed. To this end, we have to understand the basic operating principles of the PHY layer under consideration. In this paper, we use an IEEE 802.11n compliant PHY layer implementation based on [5] as a case study.

5.1 Receiver Architecture

A simplified architecture of a typical WLAN RF and PHY layer receiver is illustrated in Fig. 4. The signal

from each of the N_{RX} antennas is fed into a separate analog/RF frontend where the radio frequency signal is down-converted to baseband and digitized. After a parallel to serial converter that time multiplexes the receive chains, a digital frontend is responsible for frame start detection as well as for coarse-grained frequency synchronization. When a frame start is detected, the signals are forwarded to a time-multiplexed fast Fourier transform (FFT). The frequency-domain representation of the training sequences of the received frame are fed to a channel estimation and preprocessing module. There, the channel state information is extracted for each packet based on the training sequence and all computations for MIMO signal detection that only depend on the channel are performed. The output of the channel estimation [11, 12] and preprocessing [25, 33] module is fed into the subsequent MIMO detector together with the frequency-domain baseband samples of the data payload. The channel state information is further forwarded to the MAC layer as a basis for the RA for future data frames. After MIMO detection, the bits (with or without reliability information) are forwarded to the channel decoding module that comprises a deinterleaver and a Viterbi decoder or an LDPC decoder.

In our transmission scenario, shown in Fig. 2, the analog/RF frontend is always turned on (in receive or transmit mode) except during sleep periods. The highest work load of the subsequent digital frontend is during the frame start detection and the frequency-offset estimation in the first training phase. The FFT processes data during the training phases, the header phase, and the data payload phase of the received frame. During the training phases the channel estimation and preprocessing module is also active. The subsequent MIMO detection and the channel decoding modules process both header and data payload.

The processing rate of the different modules is indicated in Fig. 4. Most modules run at the symbol rate of the transmission that is determined by the utilized signal bandwidth of the system. The channel estimation and preprocessing module has to process data during short periods once per frame, when the channel state information is estimated. Otherwise the module is idle. The channel decoding module processes at the coded bit rate of the transmission that varies in IEEE 802.11n compliant systems over a large range from 13 Mbps to 720 Mbps, depending on the selected MCS.

5.2 LUT based Energy Model

We now consider the energy per bit of the receiver as our metric of main interest. To this end, we start by dividing the energy per received packet into three main contributions:

- A constant energy overhead for synchronization, header processing, training, channel-rate processing, and transmission of an ACK frame $e_H(\nu)$ in Joules. This part of the overall energy consumption depends partially on the choice of ν since different MCSs and different antenna configurations change also the duration of the training and have different energy costs.
- The second contribution to the overall energy consumption comprises the RF and the baseband processing during the data phase. This part is characterized by the power consumption $p_{BB}(\nu)$ in Watts and the duration of the data phase. The latter is determined by the length L and the throughput $\Phi(\nu)$ so that the corresponding contribution to the energy-per-frame amounts to $e_{BB}(\nu) = p_{BB}(\nu) \frac{L}{\Phi(\nu)}$.
- The last contribution to the total energy consumption of the PHY layer at the receiver comprises mostly the channel coding which is typically carried out on a bit-by-bit basis and can be described by the energy per bit of the channel decoding $\eta_{CC}(\nu)$ and the length of the frame L as $e_{CC}(\nu) = L\eta_{CC}(\nu)$.

Combining the three energy consumption contributions above and normalizing with the average number of successfully transmitted bits per packet, we obtain

$$\eta(\nu) = \frac{\left(\frac{e_H(\nu)}{L} + \frac{p_{BB}(\nu)}{\Phi(\nu)} + \eta_{CC}(\nu) \right)}{1 - P_e(\nu)} \quad (4)$$

as our metric of interest for optimization of the energy spend per information bit at the PHY layer of the client.

One straight forward and effective implementation for the hardware characteristics of such an energy model can be realized with a look-up table (LUT), which can be configured off-line for any target design (such as the IEEE 802.11n PHY layer that we use for the case study in this paper). For such an energy model a detailed discussion of its parameters is provided in the following subsection.

5.3 Energy Model Characterization

In order to characterize the energy model based on the IEEE 802.11n compliant PHY layer implementations presented in [5] and the analog frontend implementation discussed in [19], we will first elaborate on the components contributing to $e_H(\nu)$, $e_{BB}(\nu)$, and $e_{CC}(\nu)$. In a next step, we present an automated component wise calibration method that is based on post-layout gate-level power simulations of PHY layer implementations.

The first term of (4), the energy-overhead $e_H(\nu)$ is composed of

$$e_H(\nu) = e_{af}(\nu) + e_{df}(\nu) + e_{fft}(\nu) + e_{chpp}(\nu) + e_{det}(\nu) + e_{ack}, \quad (5)$$

where $e_{af}(\nu)$, $e_{df}(\nu)$, $e_{fft}(\nu)$, $e_{chpp}(\nu)$, and $e_{det}(\nu)$ correspond to the energy used by the analog frontend, the digital frontend, the FFT, the channel estimation and the matrix preprocessing circuit, as well as of the detector, during processing of the overhead illustrated in Fig. 2. The final component e_{ack} corresponds to the energy consumed by the PHY and RF to transmit an ACK frame to the AP.

To calibrate $e_H(\nu)$, we measured the average power consumption of the digital frontend, the FFT, the pre-processing module, the detector, and the RF. The average power consumption is then multiplied with the duration of the active time during the overhead processing of the corresponding components. While the active time during the header processing of the FFT, the preprocessing, and the detector loosely depend on the mode ν , the active time of sending an ACK frame is fixed. For a given MCS, we can assume, that the active time of each component contributing to the overhead processing energy is approximately constant.

The second term of (4), the baseband energy consumption $p_{BB}(\nu)$, is composed of

$$p_{BB}(\nu) = p_{af}(\nu) + p_{df}(\nu) + p_{fft}(\nu) + p_{det}(\nu), \quad (6)$$

where $p_{af}(\nu)$, $p_{df}(\nu)$, $p_{fft}(\nu)$, and $p_{det}(\nu)$ correspond to the power consumption of the RF, the digital frontend, the FFT, and the detector, respectively.

The last term in (4), the energy-efficiency $\eta_{CC}(\nu)$ is only composed of the channel coding module. To calibrate the LUT for that module, we measure the energy consumption per bit of the channel decoding module.

The characterization flow for the energy model is illustrated in Fig. 5. We first generate stimuli and then we monitor through simulations of the pre-synthesis hardware-description-language model the active modules for each system mode ν in order to accurately capture activity pattern of the actual hardware. The start and end time stamps of all active periods for each selected module are used to automatically generate the required scripts for post-layout power simulation. In parallel, the simulation of the post-layout circuit of the PHY layer implementation is performed. This simulation outputs value change dump (VCD) files required for accurate power estimation. In a next step, the VCD files and the generated scripts are used to generate a library storing the power values for the selected modules during the periods of interest.

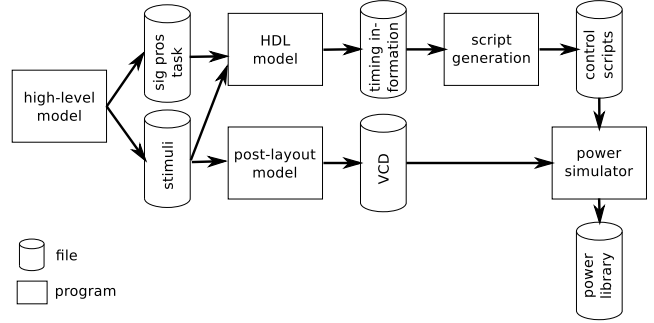


Fig. 5: Characterization of all modules in the PHY layer ASIC using an automated power library generation flow.

5.4 Discussion on the Proposed Energy Model

Before proceeding with the optimization of (4) by properly choosing ν for a large, but fixed L , a brief discussion of the implications of this model and the dependency between its variables (through the choice of ν) provides some insight into trends and ideas to simplify the estimation of the potential for energy savings.

As a starting point for our considerations, we note that in the rather generic expression in (4), $e_H(\nu)$, $p_{BB}(\nu)$, and $\eta_{CC}(\nu)$ all depend on ν . However, in practice, we note at least for large L the energy spent on overhead processing $e_H(\nu)$ is often insignificant³ compared to the other terms in (4). Furthermore, we assume that for a traditional PHY implementation, both $p_{BB}(\nu)$, and $\eta_{CC}(\nu)$ only have a limited dependency on the mode ν . Hence, choosing ν to maximize $\Phi(\nu)$ trivially minimizes the nominator of (4) with diminishing returns due to the bias terms. Unfortunately, $P_e(\nu)$ also approaches one as $\Phi(\nu)$ increases which ultimately limits the achievable goodput as well as the energy-gains. To eliminate the rather complex relationship between ν and $P_e(\nu)$, we take advantage of the presence of a fast and conservative RA. For each channel realization we divide the available modes into two groups: one that is able to get a packet across with very high probability $P_e(\nu) \approx 0$ and a second group of modes that will almost surely fail ($P_e(\nu) \approx 1$). Under these assumptions, the choice of the highest-rate mode that is still reliable is clearly the most favorite strategy and little potential would exist for further energy-efficiency optimization.

However, if different modes are able to provide significantly different $e_H(\nu)$, $p_{BB}(\nu)$, or $\eta_{CC}(\nu)$ for

³ The payload phase of the transmission for low to medium throughput modes using a large L is much longer than the duration of the header shown in Fig. 2.

the same channel realization, while still maintaining $P_e(\nu) \approx 0$, we may be able to obtain further energy-efficiency gains over the straightforward choice for the mode that provides the best goodput.

6 System Mode Extension for Enhanced Energy-Efficiency

To enable further energy-efficiency improvements beyond the ones associated with the natural choice of the highest-rate mode, the baseline PHY layer implementation must first be improved. The main idea is to first introduce a more energy proportional behavior in a sense that $e_H(\nu)$, $p_{BB}(\nu)$, and $\eta_{CC}(\nu)$ scales better with the different processing requirements of the different modes ν . In the second step, we then introduce new modes which are not necessarily optimal in terms of their goodput, but may still provide an overall energy-efficiency advantage by further reducing processing requirements at the expense of rate or reliability (error rate).

6.1 Energy Proportionality Without Impact on Goodput Through DVFS

A well designed physical layer implementation already provides a certain degree of energy-proportionality. Hence, the energy consumption of such a PHY depends more or less linearly on the number of operations required to extract the payload data from the received frame. A good example is the channel decoding, which contributes to (4) a constant energy per decoded bit. Unfortunately, such circuits do not exploit differences in the time available to process each individual bit. To take advantage of different throughput requirements in different modes, we notice that in very large integrated circuits the maximum operating frequency of a digital circuit scales approximately linearly with the supply voltage (within reasonable range). However, the power consumption of the circuit scales quadratically with the supply voltage. Therefore, we can translate relaxed throughput requirements (in terms of operations per second) into further energy-savings per operation, an idea that is commonly referred to as DVFS.

Practical applications of DVFS [20] are clearly associated with many difficulties and with overhead (e.g., voltage regulators and alike). In [37] it is shown that integrated voltage regulators with a conversion-efficiency better than 80–90% and a reasonable circuit size are feasible in modern CMOS technology. More specifically, integrated DC-DC converters providing an output current with several amps per square millimeter silicon

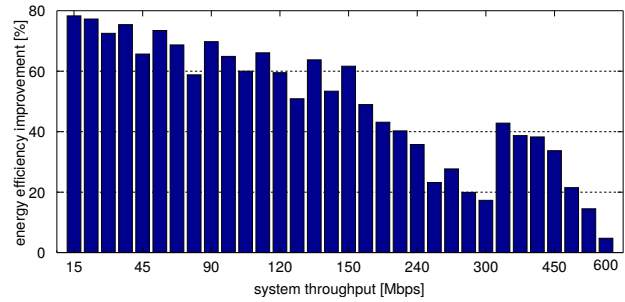


Fig. 6: Relative energy-efficiency improvement of the channel coding module using DVFS for all supported transmission schemes.

area have been proposed in [7] and [2]. Hence voltage regulators for circuits consuming few hundred milliwatt require only a small fraction of a square millimeter silicon area. However, in order to explore the limits of energy-optimal data transmission, regardless of the technical issues on circuit level we intentionally neglect the associated overhead for voltage conversion in this paper.

Application of DVFS to channel decoding: A first obvious opportunity to apply DVFS in a conventional IEEE 802.11n compliant receiver is the channel decoding module. For this module, the processing rate varies significantly (from 13 Mbps to 720 Mbps). For transmission modes with a system data rate up-to 300 Mbps, one decoding core is used. On the other hand, two cores are used for larger data rates. With DVFS, this part of the receiver can take advantage of the reduced rate and provide better overall energy-per-bit without negative impact on error rate performance of the receiver.

The relative energy savings for the channel decoding module applying DVFS are shown in Fig. 6. It can be seen, that the lower the data rate, the better the energy savings from DVFS as the clock frequency and consequently the supply voltage can be reduced. Further, it can be seen, when the number of channel decoding cores changes from one to two, as specified in the IEEE 802.11n standard, then the achievable energy-efficiency gain is again increased, as the data rate per decoder core is reduced.

6.2 Energy Proportionality With Impact on Goodput

Two other possibilities for enhancing the energy proportional behavior of the PHY layer, which can affect the goodput are the following:

Sub-optimal algorithms: A possible modification to improve energy proportional behavior is to take advantage

of the fact that in some situations different receiver algorithms with noticeable complexity difference exhibit a similar error-rate behavior. Hence, choosing an algorithm with higher computational-complexity may still not allow for a goodput improvement that out-weights the higher energy-cost.

A good example is the choice of the MIMO detector, where a solution with close to maximum likelihood (ML) or a posteriori probability (APP) performance can be combined with a low-complexity MIMO detector. The ML or APP detector provides good performance required for bad conditions, (e.g., low signal to noise ratio (SNR) or ill-conditioned MIMO channels), while the low-complexity alternative uses less energy. For conditions where the occurrence of a frame error is independent of the detector (e.g., very high SNR, very low SNR), the low-complexity alternative achieves obviously better energy-efficiency.

Exploiting active antenna selection for applying DVFS: A second opportunity to apply DVFS exists in case of reduced-complexity receiver configurations that may be suboptimal in terms of error rate performance but can be advantageous in terms of energy consumed. Since the corresponding reduced complexity modes will generally only allow for a goodput that is lower than that provided by more complex modes, they can only be advantageous when savings in $e_H(\nu)$, $p_{BB}(\nu)$, and $\eta_{CC}(\nu)$ make up for the goodput loss. The most obvious target for such additional modes is the choice of the number of active antennas N_{RX} at the receiver [28]. Such adaptation allows to roughly linearly scale the power of the RF frontend with N_{RX} , but also limits the number of spatial streams to $N_{SS} < N_{RX}$. The baseband processing can take further advantage of this scaling due to the reduced number of operations (e.g., number of multiplications and additions) but also benefits from DVFS since fewer receive chains must be processed in a time-interleaved fashion and therefore more time is available per active receive chain.

7 Evaluation

In this section we explore the limits of the gains in energy-efficiency that can be achieved with the proposed modifications to the PHY layer combined with the proposed energy-guided RA scheme. Furthermore, we compare the proposed RA schemes with the classical GG RA. We perform this evaluation on an IEEE 802.11n compliant wireless communication system for two different scenarios described in the following paragraphs.

7.1 Scenarios

In our evaluation, we consider two different scenarios: In the first scenario, the AP transmits frames with a fixed length L of 1.5 kB over a Rayleigh block-fading channel. For this simple scenario we compare the impact of all three discussed RA schemes in detail. In the second scenario the AP transmits frames with a varying length. To this end, we enable the AP to perform aggregation⁴ of up to 16 frames, each with 1.5 kB.

For both scenarios, we considered the RF given in [19] combined with the IEEE 802.11n compliant PHY layer given in [5], with up to 4 spatial streams and up to 4 receive antennas using transmissions with 40 MHz bandwidth. For channel coding we use only the convolutional code defined in the standard. The 32 mandatory MCSs defined in IEEE 802.11n result in a variable frame duration for a given L .

For 10^3 channel realizations, the reception of a frame for all considered system modes have been simulated. In particular the simulated system modes include all mandatory MCSs of the IEEE 802.11n standard. For each MCS, the number of active receive antennas varies between N_{SS} (the minimal number of receive antennas required for the specific MCS) and 4. In addition, all MCSs have been simulated with a hard-output lattice reduction aided linear minimum mean squared error (MMSE) MIMO detector (LRALD) [31], and a low-complexity soft-output MMSE detector, resulting in 112 different system modes⁵ for each MIMO detector algorithm comprising Ω_{EG} .

7.2 Ideal Estimation of Channel Conditions

As we discussed in Section 4, any RA requires an accurate estimation of $P_e(\nu)$ for the actual channel conditions. Since the purpose of this paper is to explore the potential and the limits of energy-awareness we intentionally assume a genie-aided approach for the RA. This simplification avoids introducing uncertainties due to specific RA strategies and it simplifies the full explanation of the setup under consideration. The genie provides an upper bound on goodput and energy-efficiency

⁴ Frame error rate remains independent of the number of aggregated frames (individual ACK in the same ACK frame). Still, overhead (PHY and ACK for small packets individually) can be avoided, which motivates aggregation of multiple frames.

⁵ 32 transmission modes received with 4 receive chains without applying DVFS; 32 transmission modes with 4 receive chains using DVFS at the channel decoder; 24 transmission modes using 3 receive chains and DVFS wherever applicable; 16 transmission modes using 2 receive chains and DVFS wherever applicable; 8 transmission modes using 1 receive chain and DVFS wherever applicable.

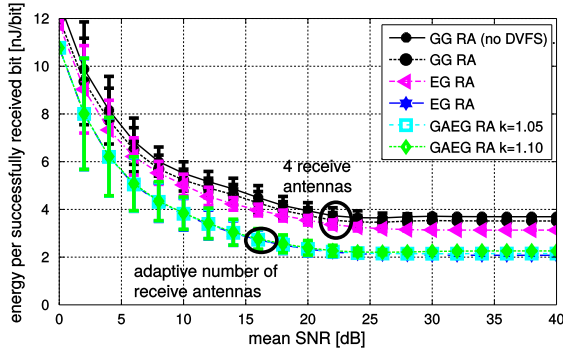


Fig. 7: Energy consumption per successfully transmitted bit.

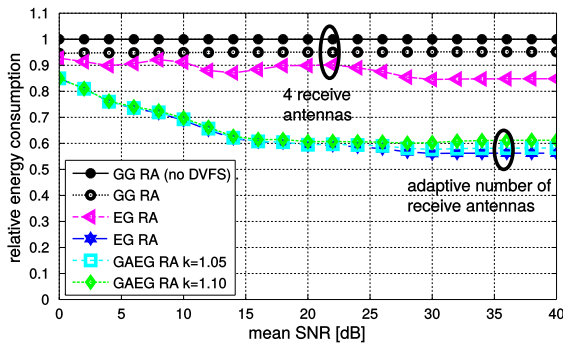


Fig. 8: Relative power consumption with respect to GG RA without DVFS.

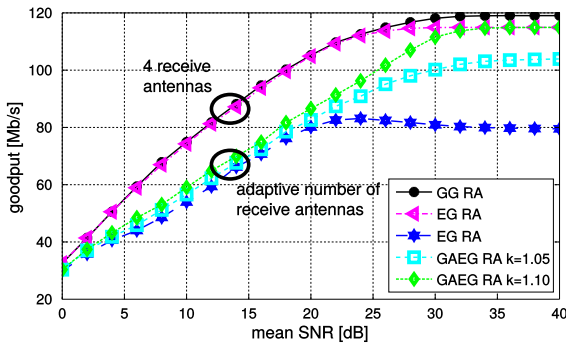


Fig. 9: Throughput achieved by the different schemes.

by perfectly predicting transmission failures by sending each packet in all modes for each channel and noise realization which – of course – is only feasible in simulation but could be approached in practice with carefully designed algorithms. Hence, $P_e(\nu)$ is either one or zero which boils down to limiting the selection to the error-free modes according to a specific optimization criterion⁶.

⁶ Practical implementation to estimate $P_e(\nu)$ can be found in [4].

7.3 Results for a Fixed Frame Length

In Fig. 7 to Fig. 9 we compare the impact of the proposed modifications in conjunction with the considered RA strategies on absolute and relative energy per bit, as well as on goodput.

We compare all obtained curves to the reference one corresponding to a GG RA without the PHY layer modifications proposed in Section 6. This baseline implementation applies no DVFS and always employs 4 receive chains resulting in the worst energy consumption per successfully received bit as shown in Fig. 7. However, such GG approach leads to the highest goodput over the entire SNR range as shown in Fig. 9.

- By assuming only PHY modifications, under a GG RA, no significant energy-efficiency gains can be achieved over the baseline implementation. In particular, GG RA never sacrifices error rate performance by reducing the number of active antennas or by using a sub-optimal MIMO detector. Hence, the only PHY modification that can be applied is DVFS in the channel coding module. Since channel decoding is only a small fraction of the overall receiver power consumption, we obtain only the marginal energy reduction of 5% as shown in Fig. 8.
- By modifying the RA to EG RA, while still applying DVFS only in the channel decoding, without the capability to scale the number of active receive chains, the energy per bit is reduced by up to 15% compared to the baseline case. Although the EG RA improves the energy-efficiency compared to the baseline implementation and the previous considered case, the restriction to exactly 4 active receive antennas does not reveal the maximum gains. Note that, the goodput does not deteriorate significantly compared to the baseline implementation.
- The maximum average efficiency gains up to 44% are obtained in case of joint application of all proposed PHY layer modifications under an EG RA (DVFS applied in the channel coding module along with DVFS based on active receive chain selection, as well as selection of an appropriate MIMO detector for the specific channel condition). However, such savings come at the cost of a reduced goodput, as shown in Fig. 9. In any case, most applications such as today's compressed HD videos on popular video-on-demand platforms or high quality audio streams require a much lower data rate than the one achieved by our approach. Hence, the proposed approach can still satisfy the data rate requirements of most popular applications while significantly improving the energy-efficiency of the battery operated device.

- For applications requiring much higher goodput and in case that minor energy-efficiency sacrifice is acceptable, then our results show that GAEG RA can provide a reasonable compromise between EG RA and GG RA. As shown in Fig. 7 and Fig. 8 GAEG RA with a factor $k = 1.05$ or $k = 1.10$ reduces the efficiency only slightly compared to the best achievable gain in case of EG RA. However, in the high SNR region (i.e., the region where high data rates can be achieved) the maximum achievable goodput is restored to 97% compared to the goodput achieved by the GG RA scheme, as shown in Fig. 9.

In order to show the efficacy of the EG RA, Tbl. 1 and Tbl. 2 demonstrate the percentaged distribution of the selected modes and the corresponding energy consumption per successfully received bit for the different SNR values. In case of 15 dB SNR, the EG RA chooses most of the time the linear MMSE MIMO detector, because the LRALD has a higher computational complexity and therefore is more power hungry, but rarely improves the reception of the frames under such channel conditions. However in the high-SNR regime, where ill-conditioned channels have a more prominent impact, the choice of LRALD is still preferable for the EG RA in many cases, as shown in Tbl. 2.

7.4 Results for Varying Frame Lengths

In order to show the relative energy improvement for the overall battery operated device in an office environment, we consider the case of aggregated frames with the IEEE TGnC [10] channel model, as described in the IEEE 802.11n standard. To this end, we evaluate for varying frame sizes in Fig. 10 to Fig. 11 the achieved energy-efficiency and goodput of the proposed PHY layer modifications under an EG RA scheme, compared with the ones achieved by the baseline implementation. For this setup, we allow the AP to aggregate up to 16 frames. Both, the GG RA and the EG RA are allowed to select the number of aggregated frames for each transmission freely such that the resulting total frame length L is between 1.5 kB and $L_{max} \leq 24$ kB, while each of the aggregated frames is acknowledged independently in a single ACK frame.

The resulting relative energy consumption with respect to the SNR and the maximum frame length L_{max} (which is an integer multiple of 1.5 kB) is depicted in Fig. 10. It can be observed that the best relative energy consumption is achieved for short frames and moderate to high SNR reaching up to 36.6%. In this moderate to high SNR region the RA can take advantage of many

Table 1: Experimental results at 15 dB SNR

GG RA			EG RA (var. ant.)			
MCS	[nJ/bit]	%	MCS	alg.	[nJ/bit]	%
12	4.21	65.0	12	MMSE	3.85	0.5
13	3.58	6.8	3	MMSE	3.27	14.6
14	3.46	0.5	4	MMSE	2.49	22.3
18	5.13	1.0	5	MMSE	2.07	4.9
19	4.39	2.4	6	MMSE	1.94	1.0
20	3.58	16.0	10	MMSE	3.44	1.5
27	3.88	8.3	11	MMSE	2.81	31.1
			12	MMSE	2.23	2.7
			18	MMSE	3.38	8.2
			19	MMSE	2.88	8.2
			26	MMSE	3.62	1.0
			27	MMSE	3.09	2.9
			14	LRLD	3.56	1.0
average: 4.05 nJ/bit			average: 2.83 nJ/bit			

Table 2: Experimental results at 30 dB SNR

GG RA			EG RA (var. ant.)			
MCS	[nJ/bit]	%	MCS	alg.	[nJ/bit]	%
22	3.07	1.0	8	MMSE	1.85	7.3
23	2.92	71.8	14	MMSE	1.93	1.0
29	2.90	0.5	15	MMSE	1.89	3.9
31	2.80	26.7	16	MMSE	1.85	7.3
			23	MMSE	1.92	0.5
			24	MMSE	1.82	48.5
			4	LRLD	1.85	12.1
			11	LRLD	1.89	1.9
			12	LRLD	1.82	17.0
			20	LRLD	1.90	0.5
average: 2.89 nJ/bit			average: 1.83 nJ/bit			

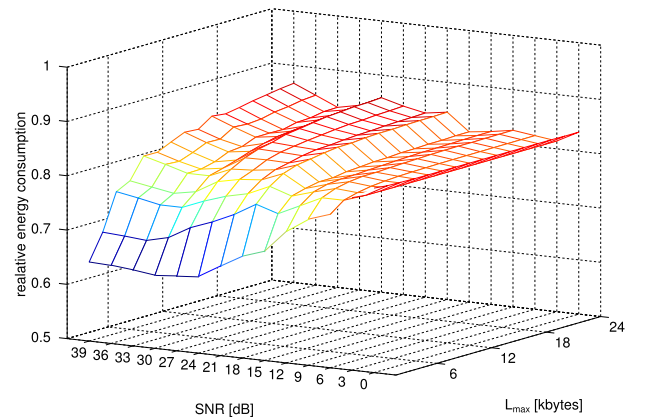


Fig. 10: Relative energy-efficiency for different maximum frame lengths and different channel quality.

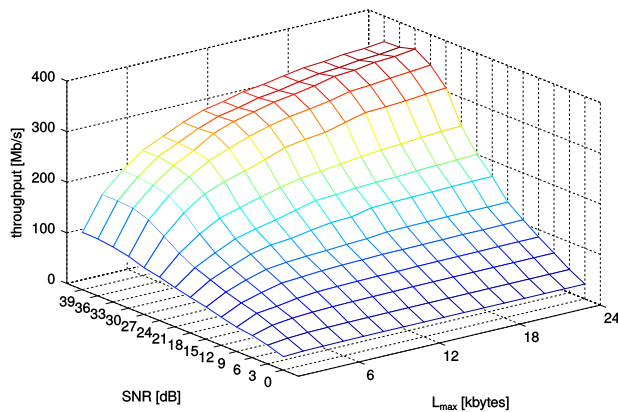


Fig. 11: Throughput of the proposed optimized system.

operation modes with a $P_e(v) \approx 0$. It is worth mentioning that the relative energy consumption of the proposed system is in all regions at least 13.3% better than that of the baseline implementation. This indicates that indeed our method can improve the energy-efficiency of mobile clients for realistic channel models.

Finally, we show the impact of the EG RA on the achieved goodput in Fig. 11. While we expected a reduced goodput for non-aggregated frames, our data shows that EG RA achieves for aggregated frames a goodput close to the goodput achievable by GG RA. Therefore there is no need for a GAEG RA and any compromise in terms of energy-efficiency can be avoided.

8 Conclusion

In this paper, we depart from conventional goodput optimized communication systems and present a cross-layer approach that involves various design, analysis and run-time techniques for improving the energy-efficiency of packet based communication systems. The proposed approach considers all blocks within the PHY layer of battery operated devices that participate in the communication link and optimizes the energy-efficiency by jointly considering the system level and hardware level parameters within the MAC and PHY layer.

The energy efficiency optimization is achieved with the introduction of an energy-guided rate adaptation scheme which is combined with several modifications of the PHY layer to enable maximum energy savings. Several new modes of operation with reduced receiver complexity are added to further improve the energy proportional behavior of the battery operated device. Fine grained energy models of PHY are developed for

propagating such behavior to the MAC layer and allowing the energy-guided rate adaptation to select the optimum system mode under any given condition and requirement.

The numerical results indicate that our approach can achieve up-to 44% energy efficiency improvement in an IEEE 802.11n system. This is achieved by properly selecting the system mode along with the right degree of voltage and frequency scaling and the selection of the appropriate algorithm for data processing when possible. We further show that a goodput-aware energy-guided rate adaptation can provide a reasonable compromise between energy and goodput if the goodput achieved by the energy-guided rate adaptation is insufficient. In any case, our work, encourages a paradigm shift towards an energy-guided RA with a properly energy proportional PHY layer, thus motivating further work towards this direction.

Acknowledgements This work was partially supported by the Hasler Foundation project WiLANCE and the EU Marie Curie DARE grant 304186.

References

1. S. Appadwedula, M. Goel, N. R. Shanbhag, D. L. Jones, and K. Ramchandran. Total system energy minimization for wireless image transmission. *J. of VLSI signal processing systems for signal, image and video technology*, 27(1-2):99–117, 2001.
2. T. Andersen, F. Krismer, J. Kolar, T. Toifl, C. Menolfi, L. Kull, T. Morf, M. Kossel, M. Brandli, P. Buchmann, P. Francese. A deep trench capacitor based 2:1 and 3:2 reconfigurable on-chip switched capacitor DC-DC converter in 32 nm SOI CMOS. In *Proc. APEC*, pages 1448–1455, Fort Worth, USA, March 2014.
3. G. Auer, V. Giannini, C. Desset, I. Godor, P. Skillermark, M. Olsson, M. A. Imran, D. Sabella, M. J. Gonzalez, O. Blume, and A. Fehske. How much energy is needed to run a wireless network? *IEEE Trans. Wireless Commun.*, 18(5):40–49, Oct. 2011.
4. S. Biaz, W. Shaoen. Rate adaptation algorithms for IEEE 802.11 networks: A survey and comparison. In *Proc. IEEE Computers and Communications*, pages 130–136, Marrakesh, Morocco, July 2008.
5. A. Burg, S. Haene, M. Borgmann, D. S. Baum, T. Thaler, F. Carbognani, S. Zwicky, L. Barbero, C. Senning, P. Greisen, T. Peter, C. Foelml, U. Schuster, P. Tejera, and A. Staudacher. A 4-stream 802.11n baseband transceiver in 0.13 μm CMOS. In *Dig. Techn. Papers, Symp. on VLSI*

- Circuits*, pages 282–283, Kyoto, Japan, June 2009. IEEE.
6. A. Carroll and G. Heiser. An analysis of power consumption in a smartphone. In *USENIX annual technical conference*, pages 271–285. IEEE, 2010.
 7. L. Chang, R. Montoye, B. Ji, A. Weger, K. Stawiasz, and R. Dennard. A Fully-Integrated Switched-Capacitor 2:1 Voltage Converter with Regulation Capability and 90% Efficiency at 2.3A/mm². In *Proc. IEEE Symp. VLSI Circuits (VLSIC)*, pages 55–56, Honolulu, HI, June 2010. IEEE.
 8. S. Cui, A. J. Goldsmith, and A. Bahai. Energy-efficiency of MIMO and cooperation MIMO techniques in sensor networks. *IEEE J. Sel. Areas in Communications*, 22(6):1089–1098, 2004.
 9. H. De Man. Ambient intelligence: Gigascale dreams and nanoscale realities. In *Proc. IEEE ISSCC*, pages 29–35, San Francisco, USA, Feb. 2005.
 10. V. Erceg, L. Schumacher, P. Kyritsi, A. Molisch, D. Baum, A. Gorokhov, C. Oestges, Q. Li, K. Yu, N. Tal, B. Dijkstra, A. Jagannatham, C. Lanzl, V. Rhodes, J. Medbo, D. Michelson, and M. Webster. TGN channel models. In *IEEE 802.11 document 03/940r4*, May 2004.
 11. S. Haene. VLSI circuits for MIMO-OFDM physical layer. *Dissertation*, Zurich, CH, May 2007.
 12. S. Haene, A. Burg, N. Felber, and W. Fichtner. OFDM channel estimation algorithm and ASIC implementation. In *Proc. ECCSC*, pages 270–275, July 2008. IEEE.
 13. D. Halperin, B. Greenstein, A. Sheth, and D. Wetherall. Demystifying 802.11n power consumption. In *Proc. Power aware computing and system*, pages 1–5, Berkeley, USA, Oct. 2010. USENIX.
 14. T. L. Jensen, S. Kant, J. Wehinger, and B. H. Fleury. Fast link adaptation for MIMO OFDM. *IEEE Trans. Vehicular Technol.*, pages 3766–3778, Oct. 2010.
 15. F. Kienle, N. Wehn, and H. Meyr. On complexity, energy- and implementation-efficiency of channel decoders. *IEEE Trans. on Communications*, 59(12):3301–3310, 2011.
 16. H. Kim, C.-B. Chae, G. de Veciana, and R. W. Heath. A cross-layer approach to energy efficiency for adaptive MIMO systems exploiting spare capacity. *IEEE Trans. Wireless Commun.*, 8(8):4264–4275, Aug. 2009.
 17. H. S. Kim and B. Daneshrad. Energy-aware link adaptation for MIMO-OFDM based wireless communication. In *Military Communications Conference, IEEE*, pages 1–7, San Diego, USA, Nov. 2008. IEEE.
 18. H. S. Kim and B. Daneshrad. Energy-constrained link adaptation for MIMO OFDM wireless communication systems. *IEEE Trans. Wireless Commun.*, 9(9):2820–2832, Sept. 2010.
 19. R. Kumar, T. Krishnaswamy, G. Rajendran, D. Sahu, A. Sivadas, M. Nandigam, S. Ganeshan, S. Datla, A. Kudari, H. Bhasin, M. Agrawal, S. Narayan, Y. Dharwekar, R. Garg, V. Edayath, T. Suseela, V. Jayaram, S. Ram, V. Murugan, A. Kumar, S. Mukherjee, N. Dixit, E. Nussbaum, J. Dror, N. Ginzburg, A. EvenChen, A. Maruani, S. Sankaran, V. Srinivasan, and V. Rentala. A fully integrated 2×2 b/g and 1×2 a-band MIMO WLAN SoC in 45nm CMOS for multi-radio IC. In *Proc. IEEE ISSCC*, pages 328–329, San Francisco, USA, Feb. 2013. IEEE.
 20. E. G. Larsson, O. Gustafsson. The Impact of Dynamic Voltage and Frequency Scaling on Multi-core DSP Algorithm Design [Exploratory DSP]. In *IEEE Signal Processing Magazine*, 28(3):127–144, May 2011.
 21. C.-Y. Li, C. Peng, S. Lu, and X. Wang. Energy-based rate adaptation for 802.11n. In *Proc. ACM Mobile Computing and Networking*, pages 341–352, Istanbul, Turkey, 2012. ACM.
 22. G. Y. Li, Z. Xu, C. Xiong, C. Yang, S. Zhang, Y. Chen, and S. Xu. Energy-efficient wireless communications: tutorial, survey, and open issues. *Wireless Communications, IEEE*, 18(6):28–35, 2011.
 23. M. Li, D. Novo, B. Bougard, C. Desset, A. Dejonghe, L. Van der Perre, and F. Catthoor. A system level algorithmic approach toward energy-aware SDR baseband implementations. In *IEEE ICC*, pages 1–6, Dresden, Germany, June 2009.
 24. M. Lin and Y. Ganjali. Power-efficient rate scheduling in wireless links using computational geometric algorithms. In *Proc. IEEE Wireless communications and mobile computing*, pages 1253–1258, Vancouver, Canada, July 2006. IEEE.
 25. P. Luethi, A. Burg, S. Haene, D. Perels, N. Felber, and W. Fichtner. VLSI Implementation of a High-Speed Iterative Sorted MMSE QR Decomposition. In *Proc. IEEE ISCAS*, 1421–1424, May 2007. IEEE.
 26. G. Martorell, F. Riera-Palou, and G. Femenias. Cross-layer fast link adaptation for MIMO-OFDM based WLANs. *Wireless Personal Communications*, pages 599–609, Feb. 2011.
 27. S. Pollin, R. Mangharam, B. Bougard, L. Van der Perre, I. Moerman, R. Rajkumar, and F. Catthoor. MEERA: Cross-layer methodology for energy efficient resource allocation in wireless networks. *IEEE*

- Trans. on Wireless Communications*, 7(1):98–109, 2008.
28. S. Sanayei and A. Nosratinia. Antenna selection in MIMO systems *IEEE Communications Magazine*, pages 68–73, Oct. 2004.
 29. C. Schurgers, O. Aberthorne, and M. B. Srivastava. Modulation scaling for energy aware communication systems. In *Proc. IEEE Low power electronics and design*, pages 96–99, Huntington Beach, California, USA, Aug. 2001. IEEE.
 30. C. Senning. Energy Efficient VLSI Circuits for MIMO-WLAN. *Dissertation*, Lausanne, CH, Oct. 2014.
 31. C. Senning, L. Bruderer, J. Hunziker, and A. Burg. A lattice reduction-aided MIMO channel equalizer in 90 nm CMOS achieving 720 Mb/s. *IEEE Trans. Circuit and System I*, 61(6):1860–1871, June 2014.
 32. C. Senning, M. Mendicute, and A. Burg. Cross layer energy-efficiency optimization for cognitive radio transceivers. In *Proc. IEEE ICASSP*, pages 3928–3932, Florence, Italy, Apr. 2014. IEEE.
 33. C. Senning, A. Staudacher, and A. Burg. Systolic-array based regularized QR-decomposition for IEEE 802.11n compliant soft-MMSE detection. In *Proc. IEEE Microelectronics*, pages 391–394, Cairo, Egypt, Dec. 2010. IEEE.
 34. M. Tauber, S. N. Bhatti, and Y. Yu. Towards energy-awareness in managing wireless LAN applications. In *Proc. IEEE Network Operations and Management*, pages 453–461, Maui, USA, Apr. 2012. IEEE.
 35. S. Thandapani and A. Kailas. An accurate energy consumption model for the physical layer in a wireless mote. In *Proc. IEEE Communications*, pages 6283–6287, Ottawa, USA, June 2012. IEEE.
 36. S.-L. Tsao and C.-H. Huang. A survey of energy efficient MAC protocols for IEEE 802.11 WLAN. *Computer Communications*, 34(1):54–67, 2011.
 37. G. Villar-Piqué, H. Jan Bergveld, and E. Alarcón. Survey and Benchmark of Fully Integrated Switching Power Converters: Switched-Capacitor Versus Inductive Approach. *IEEE Trans. Power Electronics*, 28(9):4156–4167, Jan. 2013.
 38. N. Zompakis, A. Papanikolaou, P. Raghavan, D. Soudris, and C. Francky. Enabling efficient system configurations for dynamic wireless applications using system scenarios. *J. Wireless Information Networks*, pages 140–156, June 2013.